

FDR-4 Test [8]

The FDR-4 test detects and isolates faults in the components of the shared file system configuration (local and remote FDR-4 fiber-optic very high-speed channel (VHISP) adapters, the MPX-24 VHISP channel multiplexer, and one or more SSDs).

This chapter explains the execution of the FDR-4 test. It covers the following topics:

- Understanding FDR-4
- Getting started with the FDR-4 test under UNICOS
- Execution example
- FDR-4 test menus
- FDR-4 test commands
- FDR-4 test modes
 - Local FDR loopback
 - Remote FDR loopback
 - Address loopback
 - Data error injection
 - Address/block length error injection
 - Control error injection
 - Synchronization error injection
 - Framing pulse error injection
 - Block length error injection
 - High bias error injection
 - Low lamp error injection
 - SSD loopback

8.1 Understanding FDR-4

The FDR-4 test suite was designed to diagnose the components of the shared file system configuration. Those components are: the local and remote FDR-4 fiber optic VHISP channel adapters; the MPX-24 VHISP channel multiplexer; and one or more SSDs. These components are only available on Cray Research systems with an IOS-E running UNICOS 7.C or later.

For additional information on the FDR-4 or MPX-24, see the *FDR-4/MPX-24 Engineering Note*, publication PRN-0917. (This manual is Cray Research Proprietary; dissemination of this documentation to non-Cray Research personnel requires approval from the appropriate vice president and a nondisclosure agreement. Export of technical information in this category may require a Letter of Assurance.)

The shared file system configuration is a combination of hardware and software that enables multiple Cray Research mainframes to share access to UNICOS file systems that reside entirely in the memory of several SSDs. Data passes to and from these file systems via high performance fiber-optic transmitters, receivers, and multiplexers.

When an SSD is configured into a system, it is sliced up in much the same way as a traditional disk device. Nodes in the `/dev/pdd` directory can access the entire SSD address range and are referred to as *control slices*. Nodes in the `/dev/ssdd` directory are configured to access only specific ranges of SSD addresses and are referred to as *noncontrol slices*.

Kernel support for diagnostic access to the SSD VHISP channels makes no attempt to differentiate a diagnostic request destined for FDR-4 hardware from a traditional SSD I/O operation. The only validation the kernel does is to make sure the request does not attempt to access addresses outside the valid range for the slice. However, the error injection tests are set up by turning on some high-order bits of the SSD address register.

The outcome is that, depending on the size of slice chosen, the kernel may not allow some of the error injection tests to be executed. For that reason, the error injection tests are restricted to control slices that will always have a large enough address range to accommodate SSD addresses set by OLNET.

With the exception of the SSD loopback test, the hardware forces all transfers to be four SSD blocks in size. An SSD block is sixty-four 64-bit words. For consistency, the SSD loopback test also uses this size. Also, even though the address loopback test lets you specify a block length value, the transfer will be for four SSD blocks.

8.2 Getting started with the FDR-4 test under UNICOS

To execute the FDR-4 test, you need to perform the following tasks:

1. Log in to the Cray Research system on which you intend to run OLNET.
2. Determine the names of the FDR-4 devices on your system.
3. Invoke OLNET.
4. Enter the FDR-4 test menu.
5. Set up all necessary FDR-4 test parameters.
6. Execute an FDR-4 test mode.

When entering commands in OLNET, the case of characters is important only for device names.

8.3 Execution example

This section contains an example of FDR-4 test execution. The example contains the procedure for testing the local and remote FDR-4 modules (with the local FDR loopback and remote FDR loopback test modes, respectively), the MPX-24 (with the address loopback test mode), and the entire path to the SSD (with the SSD loopback test mode).

1. On the Cray Research mainframe, enter the following command to execute OLNET:

```
/etc/diag/olnet
```

The Main menu is displayed.

2. From the Main menu, enter `FDR4` to select the FDR-4 test. The FDR-4 Test Initial menu is displayed:

```

mswan @ cool:/cool/u1/mswan/diag/dgnet

***** FDR-4 TEST INITIAL MENU *****

INITIAL MENU COMMANDS          Current Value(if applicable)
-----
DV - Device path -----> undefined
DPM - Device Path Menu

HELP - Get HELP information about this menu.
TMM - Select the OLNET FDR-4 Test Mode Menu.
RT - Return to the OLNET Main menu.

Enter a command: 
    
```

a10602

Figure 94. FDR-4 Test Initial menu

3. Select a device path with the DV or DPM command. If you enter DPM, the Device Path menu is displayed:

Device path to be tested

```

mswan @ cool:/cool/u1/mswan/diag/dgnet

116 - /dev/pdd/tst1330.3 No such device or address
117 - /dev/pdd/multi_core AVAILABLE
118 - /dev/pdd/dump AVAILABLE
119 - /dev/pdd/scr1336.2 AVAILABLE
120 - /dev/pdd/jhb AVAILABLE
121 - /dev/pdd/.Opnh_ssd AVAILABLE
122 - /dev/pdd/tk_1132.7 AVAILABLE
123 - /dev/pdd/scr1130.0 AVAILABLE
124 - /dev/pdd/tst1332.0 No such device or address
125 - /dev/pdd/tst1332.1 No such device or address
126 - /dev/pdd/tst1332.2 No such device or address
127 - /dev/pdd/tst1332.3 No such device or address
128 - /dev/pdd/ssd_blk1 AVAILABLE
129 - /dev/pdd/scr0232.0 AVAILABLE
130 - /dev/pdd/ssd_blk2 AVAILABLE
131 - /dev/pdd/ssd_blk0 No such device or address
132 - /dev/pdd/ptmp_3 AVAILABLE
133 - /dev/pdd/root_d AVAILABLE
134 - /dev/pdd/root_e AVAILABLE
135 - /dev/pdd/root_f AVAILABLE
136 - /dev/pdd/usr_d AVAILABLE
137 - /dev/pdd/usr_e AVAILABLE
138 - /dev/pdd/u23_1 AVAILABLE
139 - /dev/pdd/u45_1 AVAILABLE
    
```

a10603

Figure 95. Device Path menu

The preceding display has been scrolled down to show the device that is going to be used in this example (/dev/pdd/ssd_blk1).

4. Enter the number that corresponds to the device you wish to use, in this case, 128.

```

mswan @ cool:/usr1/mswan/hiag/dgnet
144 - /dev/pdd/root_b AVAILABLE
145 - /dev/pdd/xopen_1 AVAILABLE
146 - /dev/pdd/usr_b No such device or address
147 - /dev/pdd/usrtmp_guest AVAILABLE
148 - /dev/pdd/usr_g AVAILABLE
149 - /dev/pdd/admin_guest AVAILABLE
150 - /dev/pdd/tmp_guest AVAILABLE
151 - /dev/pdd/pnh_ssd AVAILABLE
152 - /dev/pdd/root_g No such device or address
153 - /dev/pdd/root_a AVAILABLE
154 - /dev/pdd/spool_guest AVAILABLE
155 - /dev/pdd/src_a AVAILABLE
156 - /dev/pdd/src_g AVAILABLE
157 - /dev/pdd/xopen_2 AVAILABLE
158 - /dev/pdd/src.install AVAILABLE
159 - /dev/pdd/root.install AVAILABLE
160 - /dev/pdd/usr.install AVAILABLE

Choose one of the following:
H - for Help information.
# - enter a number to select/open a device path.
<CR> - to exit this routine.

Enter a command: 128

```

a10604

Figure 96. Entering the device number

5. After a valid device path is selected, the FDR-4 Test Initial menu is updated as shown in the following display:

```
mswan @ cool:/cool/1/mswan/flag/dgnet
***** FDR-4 TEST INITIAL MENU *****

INITIAL MENU COMMANDS          Current Value(if applicable)
-----
DV - Device path -----> /dev/pdd/ssd_blk1
DPM - Device Path Menu

HELP - Get HELP information about this menu.
TMM - Select the OLNET FDR-4 Test Mode Menu.
RT - Return to the OLNET Main menu.

Enter a command: [ ]
```

a10605

Figure 97. FDR-4 Test Initial menu

6. You are now ready to test this system's local FDR-4 module. This test writes data to the local FDR-4 module and immediately reads it back and compares it.

Select the FDR-4 Test Mode menu by entering the TMM command. The FDR-4 Test Mode menu is displayed. Notice that you are told whether the chosen device represents a control or noncontrol slice, as well as the range of SSD addresses that can be accessed.

```
mswan @ cool:/usr/mswan/diag/dignet
-----
FDR-4 Test Mode Menu
Current device : /dev/pdd/ssd_blk1 (control slice)
Begins at block 02700000, for 0100000 blocks

Test Parameter Commands      Value
-----
PC - Pass count -----> 1
PT - Pattern type -----> ADDRESS
TM - Test mode -----> Local FDR Loopback

Execute & miscellaneous commands
-----
EX - Execute: Local FDR Loopback for FDR-4.
TR - FDR-4 driver trace: DISABLED
RT - Return to the Initial Menu.

Enter a command: 
```

a10606

Figure 98. FDR-4 Test Mode menu

7. Because the local FDR loopback test mode is the default, you only need to enter the EX command to cause it to execute.
8. You are now ready to test this system's remote FDR-4 module. This test writes data to the remote FDR-4 module and immediately reads it back and compares it.

To select the remote FDR loopback test mode, enter the TM command to see a list of available test modes.

```

mswan @ cool:/cool/t1/mswan/diag/dignet
The current test mode is: Local FDR Loopback.
Select one of the following or press <CR> to
leave the current test mode unchanged.

Command      Description
-----
LL  ----> Local FDR Loopback
RL  ----> Remote FDR Loopback
AL  ----> Address Loopback
DEI ----> Data Error Injection
AEI ----> Address/Block Length Error Injection
CEI ----> Control Error Injection
SEI ----> Synchronization Error Injection
FPEI ----> Framing Pulse Error Injection
BLEI ----> Block Length Error Injection
HBEI ----> High Bias Error Injection
LLEI ----> Low Lamp Error Injection

The SSD loopback test is not available when using a control slice.

Enter a command: 
    
```

a10607

Figure 99. Test Modes list

9. Enter the RL command. The FDR-4 Test Mode menu is updated to show the chosen test mode.

```

mswan @ cool:/cool/t1/mswan/diag/dignet

FDR-4 Test Mode Menu
-----
Current device : /dev/pdd/ssd_blk1 (control slice)
Begins at block 02700000, for 0100000 blocks

Test Parameter Commands      Value
-----
PC - Pass count ----> 1
PT - Pattern type ----> ADDRESS
TM - Test mode ----> Remote FDR Loopback

Execute & miscellaneous commands
-----
EX - Execute: Remote FDR Loopback for FDR-4.
TR - FDR-4 driver trace: DISABLED
RT - Return to the Initial Menu.

Enter a command: 
    
```

a10608

Figure 100. FDR-4 Test Mode menu

10. Enter the EX command.
11. You are now ready to test this system's MPX-24 module. This test reads data from the MPX module and compares it to expected data patterns. The data returned is a combination of the SSD address and block length you specify. By specifying various SSD addresses and block lengths, you can test all bit combinations for those two fields.

To select the address loopback test mode, enter the TM,AL command.

12. The FDR-4 Test Mode menu is updated to show the chosen test mode.
13. Enter the EX command.
14. Because you are currently using a control slice, you cannot execute the SSD loopback test mode. You must change to a noncontrol slice. Enter the RT command to return to the FDR-4 Test Initial menu.
15. Select a device path with the DV or DPM command. In the following display, the DV, *device name* command is used to set the device path to the one you want to use.

```

mswan @ cool:/t1/mswan/olnet
***** FDR-4 TEST INITIAL MENU *****

INITIAL MENU COMMANDS          Current Value(if applicable)
-----
DV - Device path -----> /dev/pdd/ssd_blk1
DPM - Device Path Menu

HELP - Get HELP information about this menu.
TMM - Select the OLNET FDR-4 Test Mode Menu.
RT - Return to the OLNET Main menu.

Enter a command: dv,/dev/pdd/ssd_blk2

```

a10611

Figure 101. Setting the device path

16. Enter the TMM command to get to the Test Mode menu.

```

mswan @ cool:/cool/rl/mswan/diag/dignet

FDR-4 Test Mode Menu
-----
Current device : /dev/pdd/ssd_blk2 (non-control slice)
Begins at block 02000000, for 01000 blocks

Test Parameter Commands      Value
-----
PC - Pass count -----> 1
PT - Pattern type -----> ADDRESS
TM - Test mode -----> Local FDR Loopback

Execute & miscellaneous commands
-----
EX - Execute: Local FDR Loopback for FDR-4.
TR - FDR-4 driver trace: DISABLED
RT - Return to the Initial Menu.

Enter a command: 
    
```

a10612

Figure 102. Test Mode menu

17. You are now ready to test this system's entire path from the mainframe to the SSD with the SSD loopback test mode. Data is written to, and then read from, the SSD address specified and compared.

To select the SSD loopback test mode, enter the TM,SSD command.

18. The FDR-4 Test Mode menu is updated to show the chosen test mode.
19. Enter the EX command.
20. You have completed the testing of FDR-4 components. You can either press CONTROL-d (end-of-file character) to exit OLNET or enter the commands on each screen to back out of OLNET.

8.4 FDR-4 test menus

After you initialize OLNET and access the Main menu, entering the FDR4 command displays the FDR-4 Test Initial menu.

```

mswan @ cool:/cool/t1/mswan/diag/dgnet
***** FDR-4 TEST INITIAL MENU *****

INITIAL MENU COMMANDS          Current Value(if applicable)
-----
DV - Device path -----> undefined
DPM - Device Path Menu

HELP - Get HELP information about this menu.
TMM - Select the OLNET FDR-4 Test Mode Menu.
RT - Return to the OLNET Main menu.

Enter a command: 

```

a10614

Figure 103. FDR-4 Test Initial menu

If the TMM command is entered on the FDR-4 Test Initial menu, the FDR-4 Test Mode menu is displayed.

```

mswan @ cool:/cool/t1/mswan/diag/dgnet

FDR-4 Test Mode Menu
-----
Current device : /dev/pdd/ssd_blk1 (control slice)
Begins at block 02700000, for 0100000 blocks

Test Parameter Commands      Value
-----
PC - Pass count -----> 1
PT - Pattern type -----> ADDRESS
TM - Test mode -----> Local FDR Loopback

Execute & miscellaneous commands
-----
EX - Execute: Local FDR Loopback for FDR-4.
TR - FDR-4 driver trace: DISABLED
RT - Return to the Initial Menu.

Enter a command: 

```

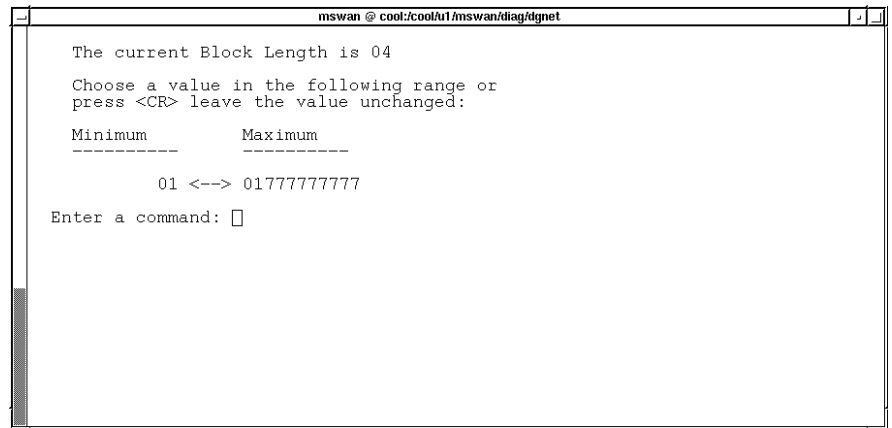
a10615

Figure 104. FDR-4 Test Mode menu

8.5 FDR-4 test commands

This section describes the commands available on the FDR-4 Test Initial menu and the FDR-4 Test Mode menu. (This section describes menu execution only. Appendix A, page 267, describes other methods of execution.) FDR-4 test commands are as follows:

<u>Command</u>	<u>Description</u>
BL	Specifies the block length to be used for the SSD loopback test mode.



a10616

Figure 105. Block length screen

BS	Specifies the bit select to be used for various error injection test modes.
----	---


```

mswan @ cool:/cool/ut/mswan/diag/dgnet
The current Bit Select is 0
Choose a value in the following range or
press <CR> leave the value unchanged:
-----
Minimum      Maximum
-----
0 <-->      7
Enter a command: █

```

a10617

Figure 106. Bit select screen

- CE** Tells OLNET to continue on error. Use the *errorfile* option to specify the file to which error output is written. These options do not appear on the FDR-4 Test menus. The CE option must be placed between the TMM and EX options in a command-line string. See Section A.2.2, page 270, for more information.
- errorfile* Specifies the file to which error output is written. This option does not appear on the FDR-4 Test menu. The *errorfile* option must be placed after the EX option in a command-line string. See Section A.2.2, page 270, for more information.
- DPM** The device path menu (DPM) command lets you display and dynamically select an SSD device path. After the DPM command has been executed, and if your system has both a `/dev/ssdd` and `/dev/pdd` directory, a menu option is displayed to choose which directory to display. If only one of these directories exists (or if you chose one of them), a menu containing paths and statuses is displayed. You can select a path from this menu.

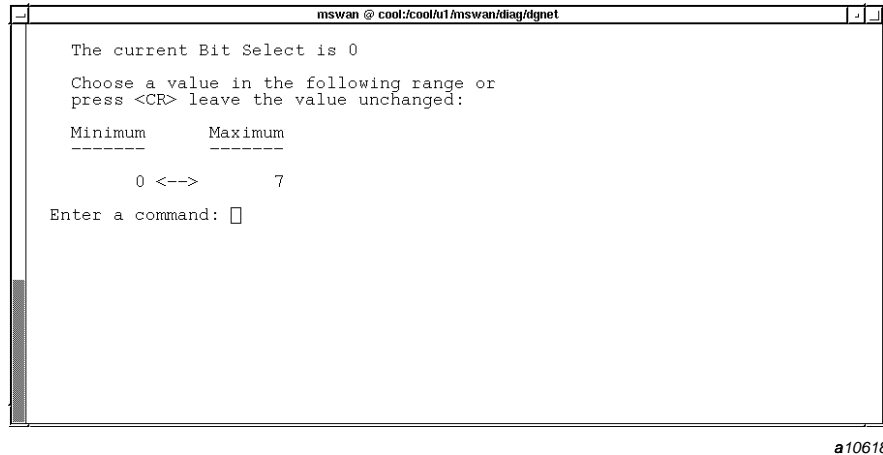


Figure 107. Paths and statuses screen

- DV Device path name. You must have read/write permission on the device path used by OLNET. Contact your system administrator to obtain these permissions.

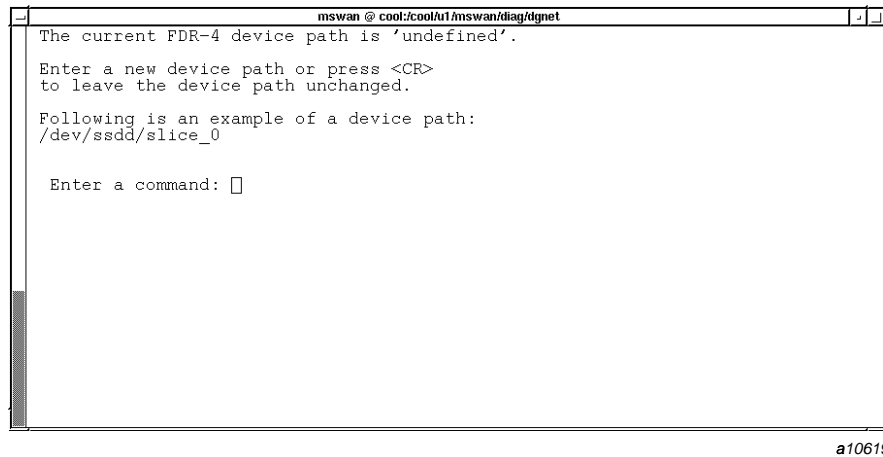
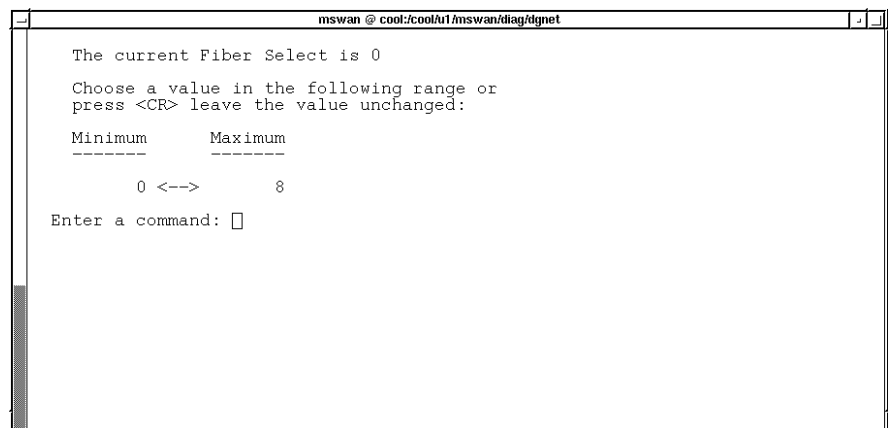


Figure 108. Device path name screen

- EX Executes the test in the test mode specified by the TM command.
- FS Specifies the fiber select to be used for various error injection test modes.



```
mswan @ cool:/u1/mswan/flag/dgnet
The current Fiber Select is 0
Choose a value in the following range or
press <CR> leave the value unchanged:
Minimum      Maximum
-----
          0 <-->          8
Enter a command: [ ]
```

a10620

Figure 109. Fiber select screen

- IL Specifies the injection location for various error injection test modes.

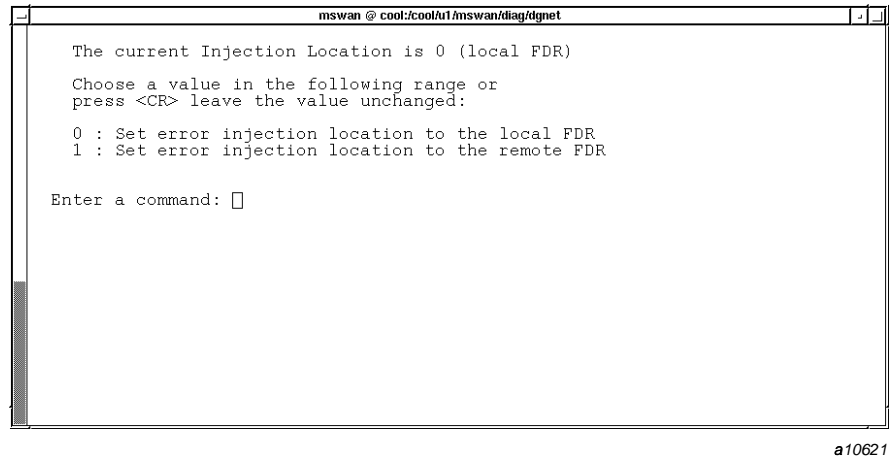


Figure 110. Injection location screen

PC, *pc* Pass count. *pc* is a value in the range 1 through 1,000,000. The default for *pc* is 1.

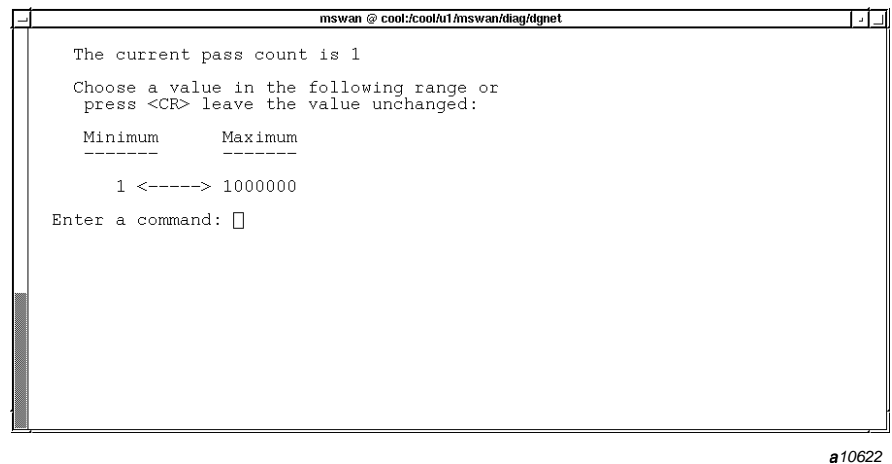


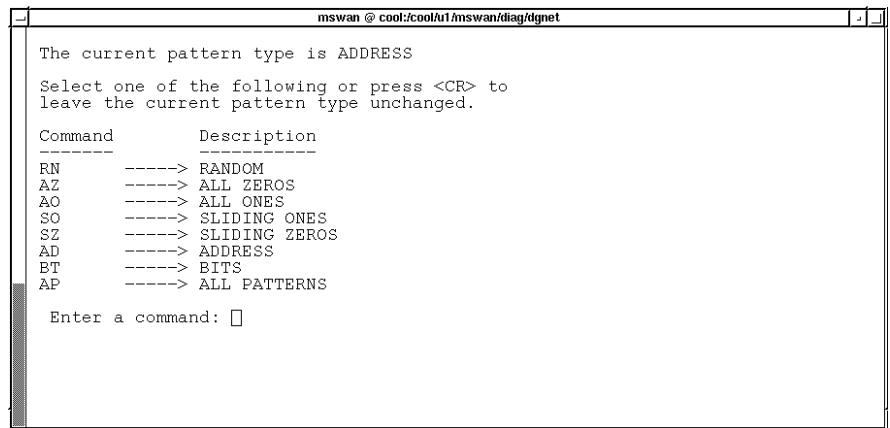
Figure 111. Pass count screen

PT, *pt* Pattern type (in 64-bit words). *pt* is one of the following values:

<u>Value</u>	<u>Pattern</u>
AD	Address (default). This sequential address pattern is incremented in each 16-bit parcel of a 64-bit word, as in the following example: 000000 000001 000002 000003 000004 000005 000006 000007
AO	All 1's.
AP	All patterns. A new pattern is generated for each message sent and received. The patterns are processed in the following order: AD, AO, AZ, SO, SZ, RN, BT.
AZ	All 0's.
BT	Bits. This pattern contains a random number of consecutive 1-bits randomly positioned within a 64-bit word, as in the following example: 000001 177770 000000 000000 000000 000000 077770 000000 177777 177777 177600 000000 000000 000000 003777 177700
RN	Random. A random pattern is generated for each message sent and received.
SO	Sliding 1's. This is a 0's data pattern in which a 1-bit is circularly shifted through each 16-bit parcel, as in the following example: 000001 000002 000004 000010 000020 000040 000100 000200
SZ	Sliding 0's. This is a 1's data pattern in which a 0-bit is circularly shifted through each 16-bit parcel, as in the following example: 177776 177775 177773 177767 177757 177737 177677 177577

The default for *pt* is AD (address pattern).

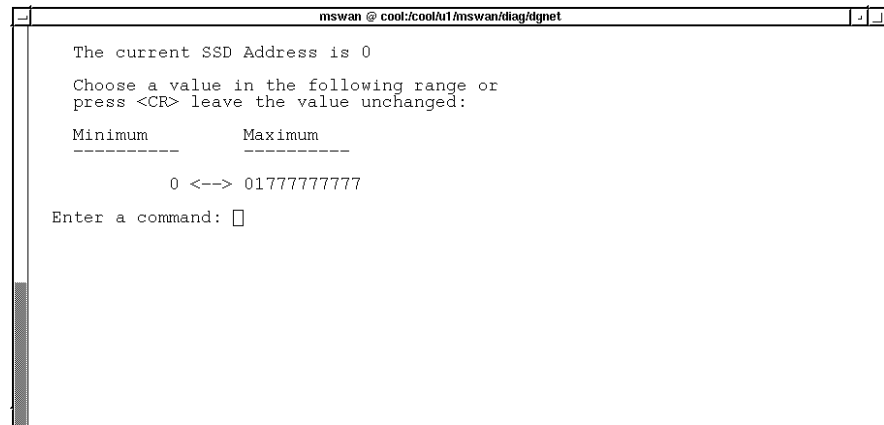
For data patterns AP, BT, and RN, OLNET builds a new pattern for each message, thereby requiring extra CPU cycles and possibly reducing the data rate (bytes/second).



a10623

Figure 112. Pattern type screen

- RT Returns to the previous menu.
- SA Specifies the SSD address for the address loopback and SSD loopback test modes.



a10624

Figure 113. SSD address screen

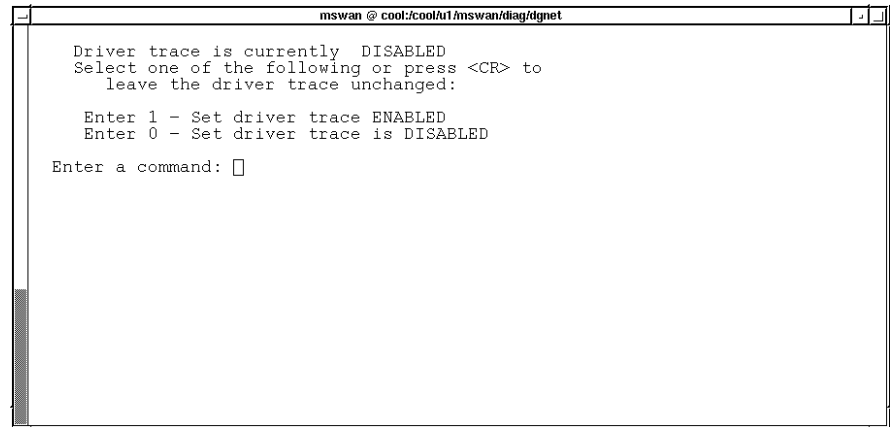
TM, <i>tm</i>	Test mode. <i>tm</i> is one of the following values:
	AEI Address/block length error injection
	AL Address loopback
	BLEI Block length error injection
	CEI Control error injection
	DEI Data error injection
	FPEI Framing pulse error injection
	HBEI High bias error injection
	LL Local FDR loopback
	LLEI Low lamp error injection
	RL Remote FDR loopback
	SEI Synchronization error injection
	SSD SSD loopback

The default for *tm* is LL (local FDR loopback).

TMM Selects the Test Mode menu.

TR Enables or disables the driver trace. The driver trace contains detailed information on the I/O system calls being executed. The

first screen displayed when you enter the TR command is as follows:



```
mswan @ cool:/u1/mswan/diag/dgnet
Driver trace is currently DISABLED
Select one of the following or press <CR> to
leave the driver trace unchanged:

Enter 1 - Set driver trace ENABLED
Enter 0 - Set driver trace is DISABLED

Enter a command: █
```

a10625

Figure 114. Driver trace screen

When you enable a trace file, the following screen is displayed:


```

mswan @ cool:/cool/t1/mswan/diag/dignet
*** SAVE TRACE INFORMATION ON A FILE ***
Enter the directory/filename.
If you want to exit this routine, just enter a <CR>.
Enter the directory/file - > 

```

a10626

Figure 115. Trace file enabled screen

8.6 FDR-4 test modes

You can execute the FDR-4 test in any of the following test modes:

<u>Test mode</u>	<u>Description</u>
Local FDR loopback	Available to either kind of slice. Data is written to, and then read from, the local FDR module and compared.
Remote FDR loopback	Available to either kind of slice. Data is written to, and then read from, the remote FDR module and compared.
Address loopback	Available to either kind of slice. Data is read from the MPX module and compared. The data returned is a combination of the SSD address and block length you specify. By specifying various SSD addresses and block lengths, you can test all bit combinations for those two fields.
Data error injection	Available only when using a control slice. You can choose which fiber and which bit on that fiber the error will be injected into as well as whether the error is injected at the local or remote FDR. The error will be detected by the MWS.

Address/block length error injection	Available only when using a control slice. You can choose which fiber and which bit on that fiber the error will be injected into. The error will be detected by the MWS.
Control error injection	Available only when using a control slice. You can choose which bit the error will be injected into as well as whether the error is injected at the local or remote FDR. The error will be detected by the MWS.
Synchronization error injection	Available only when using a control slice. You can choose which fiber the error will be injected into as well as whether the error is injected at the local or remote FDR. The error will be detected by the MWS.
Framing pulse error injection	Available only when using a control slice. You can choose which fiber the error will be injected into as well as whether the error is injected at the local or remote FDR. The error will be detected by the MWS.
Block length error injection	Available only when using a control slice. You can choose which fiber and which bit on that fiber the error will be injected into as well as whether the error is injected at the local or remote FDR. The error will be detected by the MWS.
High bias error injection	Available only when using a control slice. You can choose which fiber and which bit on that fiber the error will be injected into as well as whether the error is injected at the local or remote FDR. The error will be detected by the MWS.
Low lamp error injection	Available only when using a control slice. You can choose which fiber and which bit on that fiber the error will be injected into as well as whether the error is injected at the local or remote FDR. The error will be detected by the MWS.
SSD loopback	Available only when using a noncontrol slice. Data is written to, and then read from, the SSD address specified. The data is then compared. This is potentially the most dangerous test because SSD memory is directly manipulated.

Because of the risk of corrupting customer data, this test is not made available when using a control slice. Typically, a noncontrol slice is configured for diagnostic use.

